

Article

Using Twitter Data to Monitor Natural Disaster Social Dynamics: A Recurrent Neural Network Approach with Word Embeddings and Kernel Density Estimation

Aldo Hernandez-Suarez ¹, Gabriel Sanchez-Perez ¹, Karina Toscano-Medina ¹,
Hector Perez-Meana ¹, Jose Portillo-Portillo ¹, Victor Sanchez ²
and Luis Javier García Villalba ^{3,*}

¹ Instituto Politecnico Nacional, ESIME Culhuacan, Mexico City 04440, Mexico; ahernandezs1325@alumno.ipn.mx (A.H.-S.); gasanchezp@ipn.mx (G.S.-P.); ltoscano@ipn.mx (K.T.-M.); hmperezm@ipn.mx (H.P.-M.); jportillo@ipn.mx (J.P.-P.)

² Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK; v.f.Sanchez-Silva@warwick.ac.uk

³ Group of Analysis, Security and Systems (GASS), Department of Software Engineering and Artificial Intelligence (DISIA), Faculty of Computer Science and Engineering, Office 431, Universidad Complutense de Madrid (UCM), Calle Profesor José García Santesmases, 9, Ciudad Universitaria, 28040 Madrid, Spain

* Correspondence: javiergv@fdi.ucm.es; Tel.: +34-91-394-7638

Received: 3 March 2019; Accepted: 6 April 2019; Published: 11 April 2019



Abstract: In recent years, Online Social Networks (OSNs) have received a great deal of attention for their potential use in the spatial and temporal modeling of events owing to the information that can be extracted from these platforms. Within this context, one of the most latent applications is the monitoring of natural disasters. Vital information posted by OSN users can contribute to relief efforts during and after a catastrophe. Although it is possible to retrieve data from OSNs using embedded geographic information provided by GPS systems, this feature is disabled by default in most cases. An alternative solution is to geoparse specific locations using language models based on Named Entity Recognition (NER) techniques. In this work, a sensor that uses Twitter is proposed to monitor natural disasters. The approach is intended to sense data by detecting toponyms (named places written within the text) in tweets with event-related information, e.g., a collapsed building on a specific avenue or the location at which a person was last seen. The proposed approach is carried out by transforming tokenized tweets into word embeddings: a rich linguistic and contextual vector representation of textual corpora. Pre-labeled word embeddings are employed to train a Recurrent Neural Network variant, known as a Bidirectional Long Short-Term Memory (biLSTM) network, that is capable of dealing with sequential data by analyzing information in both directions of a word (past and future entries). Moreover, a Conditional Random Field (CRF) output layer, which aims to maximize the transition from one NER tag to another, is used to increase the classification accuracy. The resulting labeled words are joined to coherently form a toponym, which is geocoded and scored by a Kernel Density Estimation function. At the end of the process, the scored data are presented graphically to depict areas in which the majority of tweets reporting topics related to a natural disaster are concentrated. A case study on Mexico's 2017 Earthquake is presented, and the data extracted during and after the event are reported.

Keywords: twitter; data mining; word2vec; CRF; LSTM; geocoding; geoparsing

1. Introduction

Although state-of-the-art sensors can detect various natural disasters in advance (e.g., Mexico City's alarm system can timely sense earthquakes originating in the southern states) [1], the devastating consequences of these events in urban areas are usually severe. The relief efforts during and after a disaster are essential for minimizing their negative impact. These efforts are largely the result of motivating the civil society to collaborate with rescue teams, public protection agencies, and security organizations to inform, rescue, and provide restoration. The active participation of civilians in the aftermath not only strengthens the society's resiliency to a natural disaster but also improves the reliability of the information obtained from non-traditional sources [2,3]. For example, thanks to widespread wireless communication networks and mobile technologies, the dissemination of digital information now serves as a vital way to contact aid services and make appropriate decisions in a fast and more flexible manner [4]. As an example, in the 2010 earthquake in Haiti, the use of instant messages sent by civilians from different locations facilitated the reporting of trapped individuals, the provision of medical assistance, and the delivery of basic needs, such as food, water, and shelter [5]. Personal mobile phones can also be used by survivors to send messages to their relatives and the community at large about their current status, and this information can eventually be forwarded to rescue teams. Figure 1 illustrates an example of an earthquake survivor using their mobile phone to communicate with relatives.



Figure 1. An earthquake survivor uses the WhatsApp messaging system to describe their situation inside a collapsed building. The messages translated to English are *My love. The roof fell. We are trapped. My love I love you. I love you so much. We are on the 4th floor. Near the emergency staircase. There are 4 of us. My love are you ok?* As a result of these messages, rescue teams were able to save the individuals trapped in the rubble [6].

Personal mobile devices can be linked to Online Social Networks (OSNs) and enable synchronization among applications, e.g., Twitter, Facebook, and Instagram, which allows users to post and update their activities in real time [7,8]. The creation and prevalence of user-generated content [9] may include temporal and spatial information associated with different events of interest [10]. For the most part, this information is represented by georeferenced patterns that establish a relationship between the posted event and spatiotemporal characteristics of the publishing entity. As an example, an update (tweet) on Twitter that includes temporal and spatial information is shown in Figure 2.

The dynamics of OSN users and their continuous status updates, along with numerous kinds of attachments, such as photos, videos, and documents, can be considered as a social sensor because the data generated on a large scale closely resembles that acquired by traditional sensor systems [11,12]. Below are some characteristics that reinforce the notion that OSNs can be treated as social sensors [13,14]:

- **Sensor operation:** Sensors acquire data from various events as a result of observations. For example, smartphones are equipped with cameras, so users are able to obtain, process, and transmit data in real-time [12,15].
- **Processing of sensed data:** When the information acquired by traditional sensing systems is processed, geographic information is available if navigation systems, e.g., GPS, are used. Information posted on OSNs may include either specific locations or textual descriptions of a place during an event. Moreover, users can reply, comment, and retransmit an update [16].



Figure 2. A tweet providing the location (spatial information) of a collapsed building, along with a timestamp (temporal information), one day after the 2017 earthquake in Mexico City. The message translated to English is: *Mexico. Preliminary damage report #Earthquake in #CdMx Zapata and Peten and Division del Norte collapsed building...* It is worth noticing that some users mention places using hashtags. In this example a hashtag #CdMx was used to refer to Mexico City.

Twitter has been popularized for the ease of reading, writing, and collecting data, which are published on a constant basis. Twitter allows users to publish opinions, sentiments, and observations, as well as update their statuses in an asymmetrical form (unlike other OSNs, such as Facebook, a Twitter user's newsfeed, mentions, and replies remain public by default). Recently, Twitter has been the center of attention in different research fields related to Marketing, Social Sciences, Natural Language Processing (NLP), Opinion Mining, and Predictive analysis [17]. Additionally, several applications are being developed to analyze Twitter data related to daily-life matters. For example, during electoral events, the work in [18] confirmed that a high rate of tweets posted by users shows a correlation with the performance of candidates and the public's preferences. Event prediction and monitoring can then be carried out by applying connective action theory that links a live event with the reactions of users [19]. For example, it has been demonstrated that events with a negative impact on society can motivate hacker activists to perpetrate cyber attacks [20]. Twitter can then be used as an alternative engine for exchanging information related to natural disasters, such as fires, floods, hurricanes, and earthquakes. Moreover, recent research has demonstrated [21] that Twitter can also be a source of information for spreading awareness of ecological phenomena with well-defined temporal patterns.

In this work, a methodology is proposed that uses Twitter as a social sensor for natural disasters by exploiting the spatial and temporal information associated with the observations and experiences posted by users. The aim of our social sensor is to provide useful geo-temporal patterns that may appear during and after the occurrence of an event, which may be useful to assess the extent of the damages.

By default, tweets are short messages of a maximum of 280 characters in length. Tweets can include well-defined geographic data provided by GPS or manual check-ins. However, it has been reported that only a very small percentage of Twitter users use navigation systems or register places to reference their status [22]. Given this difficulty in determining location, some studies have proposed estimating

the location of a tweet by exploiting some of Twitter's available features, including searching for updates related to certain events within a known geographical region [23], grouping textual patterns associated with user language [24], and parsing Twitter geo-objects to calculate the approximate coordinates from statuses that depict well-defined places [25]. Further, to tackle these limitations, the textual content of a tweet can be examined to determine whether a location is mentioned.

An important contribution of our work is to expand on the idea of examining the textual content of tweets by inspecting the so-called *toponyms* (places implicitly described in a text) from the surge of tweets that emerge during and after a natural disaster. To this end, our proposed approach employs Named Entity Recognition (NER), which is an information extraction method for finding and sorting named entities into pre-defined tags (persons, locations, and organizations) [26,27]. This is achieved by breaking down tweets into word units and classifying them into named entity tags so that a *toponym* can be discovered and geocoded (estimating its spatial information in terms of latitude and longitude coordinates). Detecting places is not a trivial task, and major challenges associated with tweets must be addressed, such as the ungrammatical nature of tweets, as well as informal abbreviations and lexicons (for example, mentioning a location using a hashtag). With respect to temporal information, we cluster values of the time and duration of tweets connected to the event of interest by similarity within a window of time [28]. To capture the semantic, morphological, and contextual richness of each word in a tweet, we perform a word-level analysis by using Word Embeddings [29,30], a widely used algorithm that transforms similar words into a continuous vector space. A sentence-level analysis is subsequently performed to extract semantic and syntactic information from each tweet by employing a Bidirectional Long Short-Term Memory (biLSTM) network [31,32], which is capable of using long-ranged symmetric sequence contexts. After training a Conditional Random Field (CRF) classifier [33] with biLSTM output sequences and their corresponding NER target classes, our methodology predicts locations from tweets. Finally, it applies a Kernel Density Estimation (KDE) algorithm [34] to the classified locations to compute various hotspot heat maps for the event of interest.

We have tested the proposed sensor with (Spanish) tweets from the 2017 Mexico City earthquake. Based on our evaluations, our sensor can accurately capture information that can help authorities, institutions, and volunteers to detect major risk areas and locate missing individuals and shelters.

2. Related Work

The detection of events related to natural disasters using OSN data has been the subject of recent research in the fields of sensors, natural language processing, and automatic and statistical learning. The common goal is to detect, monitor, and disseminate information about the event in a timely manner with some degree of trust. As described in [35], Twitter has been recently used as a platform to post diverse information related to various natural disasters, such as wildfires [36], floods [35], hurricanes [37], and earthquakes [38], and it has resulted in situational awareness. Table 1 summarizes the contributions of important works that employed data extracted from Twitter and other OSNs to sense natural disasters.

Table 1. Related work that contributes to natural disaster sensing using data extracted from Twitter and other OSNs.

Title	Description
Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors.	To detect a target event, this work classifies tweets on the basis of features such as keywords and the number of words given a context. Then, the methodology estimates a probabilistic spatiotemporal model to find the center and the trajectory of the target event. To this end, each Twitter user is assumed to be a sensor. Then, Kalman particle filtering is applied for location estimation with ubiquitous/pervasive computing. The authors claim that a 96% probability of correctly detecting an earthquake can be achieved by monitoring textual features [39].

Table 1. Cont.

Title	Description
Public health implications of social media use during natural disasters, environmental disasters, and other environmental concerns.	This work analyzes how social media can be used to disseminate information, predict data, and provide early warnings within the context of environmental awareness and health promotion. The work also analyzes how social media can be used as an indicator of public participation in environmental issues. The authors found evidence supporting social media as a useful surveillance tool during natural disasters, environmental disasters, and other environmental concerns. The work shows that public health officials can use social media to gain insight into public opinions and perceptions. Moreover, the work shows that social media allows public health workers and emergency responders to act more quickly and efficiently during crises [40].
Real-Time Crisis Mapping of Natural Disasters Using Social Media.	In this work, the authors propose a social media crisis mapping platform for natural disasters that uses statistical analysis with geoparsed real-time tweet data streams matched to locations from gazetteers, street maps, and volunteered geographic information. Geoparsing results are benchmarked against existing published work and evaluated across multilingual datasets. Two case studies are presented to compare five-day tweet crisis maps compiled from verified satellite and aerial imagery sources for official post-event impact assessment by the US National Geospatial Agency [41].
Tweedr: Mining Twitter to Inform Disaster Response.	In this paper, the authors introduce Tweedr, a Twitter-mining tool that extracts actionable information for disaster relief workers during natural disasters. The Tweedr pipeline consists of three main parts: classification, clustering, and extraction. In the classification phase, they use classification methods, namely, Latent Dirichlet Allocation (LDA), Support Vector Machines (SVM), and Logistic Regression, to identify tweets reporting damage or casualties. In the clustering phase, they use filters to merge tweets that are similar. Finally, in the extraction phase, they extract tokens and phrases that report specific information about different classes of infrastructure damage, the types of damage, and casualties [42].
A Linguistically-driven Approach to Cross-event Damage Assessment of Natural Disasters from Social Media Messages.	In this work, the authors focus on the analysis of Italian social media messages for disaster management. Their aim is to detect those messages conveying critical information for the damage assessment task. The main novelty of this study is the focus on out-of-domain and cross-event damage detection and the investigation of the most relevant tweet-derived features for these tasks. They conducted different experiments by resorting to a wide set of linguistic features to qualify the lexical and grammatical structure of a text, as well as ad-hoc features specifically extracted for this task [43].
Combining Machine Learning Topic Models and Spatio-temporal Analysis of Social Media data for Disaster Footprint and Damage Assessment.	The authors propose a crisis mapping system by analyzing the textual content of disaster reports from a twofold perspective. A damage detection component employs an SVM classifier to detect mentions of damage among emergency reports. A novel geoparsing technique is proposed and used to perform message geolocation. They report a case study to show how the information extracted through damage detection and message geolocation can be combined to produce accurate crisis maps. The crisis maps detect both highly and lightly damaged areas, thus opening up the possibility to prioritize rescue efforts where they are most needed [44].
From Social Sensor Data to Collective Human Behaviour Patterns: Analysing and Visualising Spatio-temporal Dynamics in Urban Environments.	This paper presents an approach to analyzing social media posts to assess the footprint of and the damage caused by natural disasters by combining machine learning techniques (LDA) for semantic information extraction with spatial and temporal analysis (local spatial autocorrelation) for hotspot detection. The results demonstrate that earthquake footprints can be reliably and accurately identified. The results also show that a number of relevant semantic topics can be automatically identified without a priori knowledge, revealing clearly differing temporal and spatial signatures. Furthermore, a damage map that indicates where significant losses have occurred is also presented [11].
The Performance of Publicness in Social Media: tracing patterns in tweets after a disaster	The authors propose a computer-assisted discourse analysis—specifically, a corpus-linguistic-informed analysis of half a million tweets—in order to describe four main public discursive moves that were prevalent during the earthquake in Aotearoa, New Zealand, in 2011. The final results describe how people employ their social media communication at critical, reflexive moments, such as in the aftermath of disaster [45].
Spatio-Temporal Distribution of Negative Emotions in New York City After a Natural Disaster as Seen in Social Media	In this paper, the authors propose a sentiment analysis technique termed <i>Extracting the Meaning Of Terse Information in a Visualization of Emotion (EMOTIVE)</i> , which uses spatial regimes regression to find significant associations of negative emotional responses by using social media posts over space and time in the aftermath of a natural disaster. The process can be used as a guide to identify those areas and populations in the most need of care [46].

Although the state of the art provides significant advances in geoparsing *toponyms* from OSNs using NER techniques [39–44], some challenges still exist. Important challenges are described below:

- *Vector space feature representations*: a vector space model can capture the relevance of words by assigning them a numerical weight; then, each sentence can be represented as a sparse or dense vector of a vocabulary of size V . Some algorithms include One-hot-Encoding, Bag-of-Words, and Tf-IDF (Term frequency–Inverse document frequency) [47]. Such type of codification may fail to preserve semantic, syntactic, and linguistic features, as it cannot establish relationships and similarity patterns among words in a given corpora, making it difficult to examine transitions between contiguous data.
- *Algorithm selection*: Geoparsing techniques based on NER require a suitable algorithm with minor preprocessing to train sequential structures such as tweets; more specifically, long contextual information should be considered in both directions of a word of interest. For this reason, approaches that employ SVM, Feed-Forward Neural Networks, Decision Trees, and single CRF classifiers may be unsuccessful as they assume that words are independent of each other, and rely on previous feature extraction steps. Recent approaches based on feed-forward algorithms for NER classification may have several disadvantages, as tabulated in Table 2.

Table 2. Disadvantages of algorithms employed for NER classification

Algorithm	Disadvantages for NER Tasks
Decision Trees (DT) and Random Forests (RF)	In [48], the authors conclude that DT and RF can create useful rules for sentence segmentation and partial parsing for NER classification and <i>toponym</i> identification, but they do not adequately consider linguistic or semantic knowledge. DT and RF are, unfortunately, prone to overfitting, and their complexity may be exponential in online learning scenarios and for high-dimensional sets, such as textual corpora.
Naive Bayes (NB)	The authors of [49] suggest that NB can be used for nonlinear NER classification by computing the posterior probability of a word associated with a specific tag. Despite this, NB approaches may fail if they do not consider an intermediate representation of the word-by-word composition of each sentence, especially for examining the sequential relationships among the input words.
Support Vector Machines (SVM)	SVM-based applications have been widely used for NER tasks [50] to efficiently increase accuracy scores. Although designed to maximize the decision boundaries for binary classification problems, <i>kernel tricks</i> can help to adapt nonlinear data to different dimensions, as well as adjust training steps for multi-class datasets. In any case, SVM architectures depend on previous handcrafted feature extraction techniques, which may result into high-dimensional and sparse results, making it time-consuming for sequential problems such as NER.
Single Conditional Random Fields (CRF)	CRF is one of the top-ranked generative algorithms used for NER, as studied in ref. [50]. Its main difference from discrete classification is that data are represented as sequences of tags with mapped classes for each one. Predictions are presented by maximizing the log-likelihood of the state sequences given the observed classes. Indeed, similar to SVM and DT, generative models such as CRF applied on its own cannot properly generalize long-range dependencies, such as the contextual and lexical features needed for NER.

In order to clearly define each step of the proposed methodology, the following research questions are raised:

1. How can the semantic, morphological, and linguistic textual patterns that properly represent a word and its surrounding context be preserved?
2. How can a sequential labeling problem such as NER be addressed by capturing contextual information in both directions of a word of interest, and how can it be classified as a *toponym*?
3. Why is it important to scrutinize neighboring named entity tags as state sequences at a sentence level?
4. How should geocoded data and clustered temporal information be statistically scored to depict the dynamics during and after an event?

This work aims to answer these questions. Our main contributions are summarized as follows: (1) A text preprocessing module to remove noisy textual features; (2) Word embedding representations to depict each word of interest, keeping the semantic, syntactic, and linguistic relevance; (3) An NER-based geoparsing strategy (*toponym extractor*) based on a Recurrent Neural Network (RNN) with a CRF output layer to determine the word embeddings that form a tweet and their mapped states (named entity tags); (4) A Geocoder to query Google Maps API with each *toponym*, thus presenting results in latitude and longitude values; (5) A KDE algorithm to graphically depict hotspots from clusters of geocoded *toponyms* in the same spatial area during and after the event of interest.

3. Proposed Methodology

The block diagram of the proposed sensor is depicted in Figure 3. Each block is briefly described next:

Training data

1. *Training set and Named Entity tags*: a training set is prepared with tokenized (segmented text into word units) sentences and manually inspected tweets, along with their corresponding NER tags (Named Entity classes).
2. *Preprocessing*: a step aimed to clean data by removing noisy information, e.g., unnecessary punctuation marks mistakenly added to words, extra spaces, extra line breaks, and bad character encodings, such as emoticons or emojis.
3. *Word embeddings*: Word2Vec [29,30], a well-known word embedding learning algorithm, is used to transform the preprocessed tokens into an n -dimensional word vector representation of neighboring context similarity.
4. *biLSTM and CRF*: biLSTM [31,32] is an RNN variation with extended memory capabilities. In this step, word embeddings are used for training by examining words in both directions. This is achieved by adding two separate hidden layers to provide past and future contextual information in specific time frames. Finally, a CRF output layer [33] is used with biLSTM output sequences to exploit their inherent neighboring entity tag transition states over the whole tweet.

Sensing stage

1. *Twitter data*: Tweets are scraped using a tool developed in [51]. To be able to filter a meaningful portion of tweets, a compound of queries containing information depicting urban spaces, words, and hashtags related to a natural disaster are stored and grouped into one of the following topics: $T \in \{\text{disaster areas, missing individuals, shelters}\}$.
2. *Preprocessing and Word Embeddings*: Tweets scraped in real time are cleaned and transformed into their word embedding representations using the same process as that used in the Training Stage.
3. *Classification model*: This model is obtained during the training stage and comprises a generalization of word embeddings and entity tags to be used to classify incoming tokenized tweets into named entity tags.
4. *Geoparsing and Geocoding*: Classified tokens are presented as single and sentential words that must be joined correctly to form a *toponym*. A geocoder is developed to resolve *toponyms* to their geographical coordinates by querying Google Maps [52] API and obtaining spatial information in terms of real latitude and longitude values.
5. *KDE*: The occurrence of geocoded *toponyms* in the same spatial region represents the event dynamics as it is a means of understanding what, when, and where users are reporting during and after the event. Such occurrence can be graphically analyzed by using KDE, an algorithm capable of estimating the density of reported locations within some topic $\in T$ occurring in a well-defined space, such as a hotspot heat map.

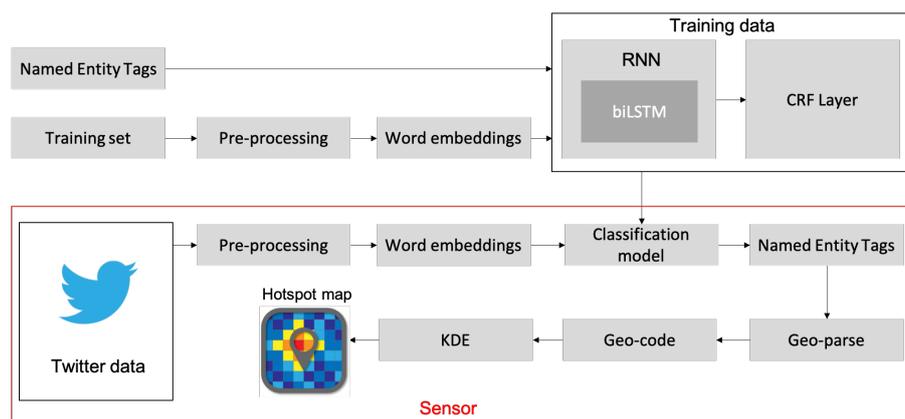


Figure 3. Proposed Twitter-based social sensor for natural disasters.

3.1. Named Entity Tags

Named entities are sequences of words that denote names of things, such as proper names, streets, avenues, and organizations [53,54]. A named entity tag is a discrete class that describes the entity type. Table 3 lists the set of named entities and tags used in the proposed sensor.

Table 3. Entity tags used for classification.

Named Entity Tag Type	Description
LOC	Location representation, e.g., a street, avenue, region, or country
ORG	Reference to an organization, institution, or establishment
PER	Reference to a person or a group of people
O	Any other criteria

3.2. Training Set

To build an NER training set, the CoNLL-2002 [55] Spanish dataset was merged with manually inspected tweets using terms in Mexican Spanish related to natural disasters. Tweets were collected by exploring historical messages and hashtags related to Mexico City's major earthquakes on the following dates: 8 September 2017; 7 June 2014; 17 April 2014; and 20 March 2012. Each training sample comprises a word, w_i , and its corresponding named entity tag, y_i . An empty entry in the training set, X , represents a sentence boundary. The CoNLL-2002 dataset contains named entity tags with a prefix indicating their position in the sentence, for example, I-LOC indicates that the position is *inside* the sentence and B-LOC indicates that the position is at the beginning of the sentence; thus, we mapped the CoNLL-2002 tags to generic tags, as listed in Table 4.

Table 4. Named entity tags used in the training set.

CoNLL-2002 Tag	Generic Tag
I-LOC, B-LOC	LOC
I-ORG, B-ORG	ORG
I-PER, B-PER	PER
O, I-MISC, B-MISC	O

To illustrate how manually inspected samples (in Spanish) are added to the training set, X , Table 5 lists some example tweets whose constituent words are assigned to a named entity tag.

Table 5. Examples of tweets with their corresponding named entity tags.

Tweet in Spanish	English Translation
$\overset{O}{\text{ayuda}} \overset{O}{\text{gente}} \overset{O}{\text{atrapada}} \overset{O}{\text{en}} \overset{LOC}{\text{edificio}} \overset{LOC}{\text{Alvaro}} \overset{LOC}{\text{Obregon}}$	help people trapped in a building located in Alvaro Obregon
$\overset{O}{\text{derrumbe}} \overset{O}{\text{en}} \overset{O}{\text{el}} \overset{LOC}{\text{multifamiliar}} \overset{O}{\text{de}} \overset{LOC}{\text{av.}} \overset{LOC}{\text{tlalpan}}$	a collapsed department building on tlalpan avenue
$\overset{O}{\text{se}} \overset{O}{\text{cayó}} \overset{O}{\text{el}} \overset{ORG}{\text{Soriana}} \overset{O}{\text{de}} \overset{LOC}{\text{Taxqueña}}$	Taxqueña's Soriana has fallen down

A total of 312,138 different words were used as inputs to a word embedding transform function, as described in Section 3.3.

3.3. Word Embeddings

Word-level analysis, also known as word embedding, is a widely used language model transformation [29,30,56,57] whose purpose is to describe words within a certain context. Each word is mapped to a new representation on the basis of its neighboring word co-occurrences in view of semantic, morphological, and linguistic patterns. The main advantage of this kind of language model transformation is its lexical richness, which makes it suitable for handling the non-grammatical nature of data extracted from OSNs; in other vector representation models, this can result in high dimensional data, thus bad weighting factors.

We show the advantages of word embeddings by taking a text describing a location, *Avenida Alvaro Obregon # 286* (a location entity in Mexico City), which can be written in different ways: *Av alv Obregon num 286*, *Ave Alvaro Obregon # 286*, *av. Alvaro Obregon 286*, or *Alv. Obregon 286*. Such variants could be a serious challenge if a feature extraction method that relies on normalizing the frequency of words contained in a document set is employed, e.g., a Vector Representation Model such as the Term frequency–Inverse document frequency (Tf–Idf) algorithm [58]. The generalization employed by such methods may imply a high-dimensional set with a complex interpretation. Instead of using weighting factors, tweets are transformed into vector representations using the Word2Vec-Skip-Gram model [29,30,59]. The Skip-Gram model is widely used for NLP-related tasks by transforming the words composing a sentence into n -dimensional vector representations given a desired context, w_ψ . The model then computes the conditional probability, $p(w_\psi|w)$, of a word, w , from a given corpus of tweets, X . A series of iterations must be performed to tune a parameter β that maximizes the probability over X , as formulated in Equation (1):

$$\operatorname{argmax}_{\beta} \prod_{w \in X} \left[\prod_{w_\psi \in \Psi(w)} p(w_\psi|w; \beta) \right], \quad (1)$$

where $\Psi(w)$ is a set of contexts describing a word w . To parameterize the Skip-Gram model, it is necessary to make use of the conditional probability $p(w_\psi|w; \beta)$ through a Softmax function, as described in Equation (2):

$$p(w_\psi|w; \beta) = \frac{e^{\mathbf{v}_{w_\psi} \cdot \mathbf{v}_w}}{\sum e^{\mathbf{v}_{w'_\psi} \cdot \mathbf{v}_w}}, \quad (2)$$

where $\mathbf{v}_w \in \mathbb{R}^n$ and $\mathbf{v}_{w'} \in \mathbb{R}^n$ are the input and output vector representations, respectively, of a word w .

3.4. Bidirectional Long Short-Term Memory Network

Long Short-Term Memory (LSTM) networks are variants of RNNs and used to solve a wide range of sequential data problems, such as Sentimental Analysis, Speech Recognition, and NER

applications [32,60], since they have the ability to capture and exploit historical and long-range dependencies with variable lengths, for example, by capturing past (from the previous words) and future (from the next words) information of a word in a tweet. In text-processing tasks, LSTM networks take words as inputs in a distributed representation of n -dimensional vectors with continuous values, in which each word belongs to a finite vocabulary $V \in \mathbb{R}^{n \times V}$. In this work, the inputs are the word embedding representations, \mathbf{v}_w , previously transformed by the Skip-Gram model. An LSTM network is constructed with hidden layer updates built into a memory cell, c . Each memory block is connected recurrently with an input, forget, and output gate, represented by i , f , and o , respectively. When trained, these gates are able to write, read, and reset information. In Equation (7), each gate is defined:

$$i_t = \sigma(W_{x_i}x_t + W_{h_i}h_{t-1} + W_{c_i}c_{t-1} + W_{0,i}), \quad (3)$$

$$f_t = \sigma(W_{x_f}x_t + W_{h_f}h_{t-1} + W_{c_f}c_{t-1} + W_{0,f}), \quad (4)$$

$$c_t = f_t c_{t-1} + i_t \tanh(W_{x_c}x_t + W_{h_c}h_{t-1} + W_{0,c}), \quad (5)$$

$$o_t = \sigma(W_{x_o}x_t + W_{h_o}h_{t-1} + W_{c_o}c_{t-1} + W_{0,o}), \quad (6)$$

$$h_t = o_t \tanh(c_t), \quad (7)$$

where σ is the sigmoid function; i_t , f_t , and o_t are the outputs of the input, forget, and output gates, respectively; c_t is the output of the cell gate constrained to the size of the hidden vector, h_t ; and W and W_0 are the weights and bias vectors, respectively.

Although RNNs, including LSTM networks, are useful for working with sequence tagging, they may fail if only past contexts (previous words) are considered. In order to account for the subsequent context, two extra hidden layers are included to process data in a bidirectional fashion. This adaptation is known as a Bidirectional Long Short-Term Memory (biLSTM) network. By training a biLSTM network, the predictive capabilities of a CRF output layer are enhanced by taking advantage of historical information from past vector representations (via forward states) and future vector representations (via backward states). In order to illustrate how a biLSTM works, an example is shown in Figure 4.

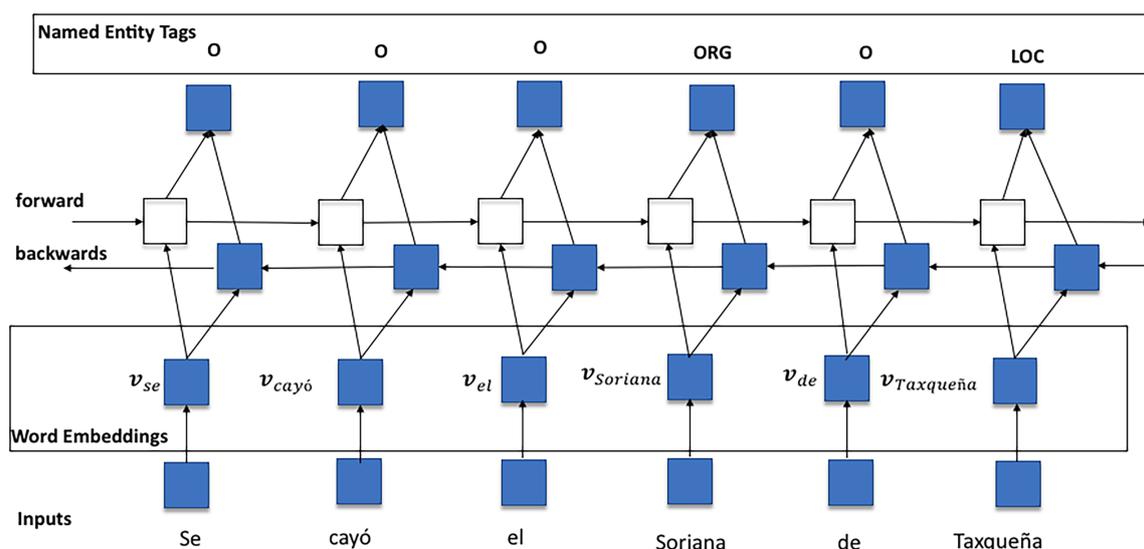


Figure 4. A biLSTM network for NER tasks. English Translation: *Taxqueña's Soriana has fallen down.*

3.5. Conditional Random Fields

Conditional Random Fields (CRF) [33] are one of the most widely used generative classifiers intended to address NER tasks [61–63] as long as their focus is on sequential data. To predict named

entity tags, a word-level examination is conducted with a set of sorted and sequential words mapped with an internal state of transitions produced by their corresponding entity tags. When combined with biLSTM networks, the resulting architecture can efficiently process NER sequences with past and future word embedding representations and efficiently predict the entity tag. To this end, a matrix of scores must be computed from the biLSTM outputs, denoted by $f_{\theta}([\mathbf{v}_{\mathbf{w}}]_1^T)$, in which $[\mathbf{v}_{\mathbf{w}}]_1^T$ is a sequence of word embeddings associated with a parameter θ , which denotes the score of the i -th named entity tag and the t -th word embedding. A transition score, $[A]_{i,j}$, is defined to shape the variation from the i -th state to the j -th state in each pair of consecutive time steps. Lastly, to score a sequence of word embeddings, $[\mathbf{v}_{\mathbf{w}}]_1^T$, with a path of tags, $[y_{i,k}]_1^T$, the sum of the total scores and network scores is calculated according to Equation (8):

$$\sum_{t=1}^T ([A]_{[y_i]_{t-1}, [y_i]_t} + [f_{\theta}]_{[y_i]_t, t}). \quad (8)$$

Algorithm 1 depicts the steps taken to train the biLSTM-CRF network; *batch* denotes the number of sequences of word embeddings, *epochs* indicates the number of epochs used for training, and $[A]_{i,j}, \theta$ are the parameters to update.

Algorithm 1: Training Samples.

```

function BILSTM-CRF( $\theta, batch, epochs$ );
  foreach epoch do
    foreach batch do
      1. compute the bidirectional LSTM-CRF forward pass for  $[\mathbf{v}_{\mathbf{w}}]_1^T$ 
          • forward pass for forward LSTM state, evaluating  $f_{\theta}([\mathbf{v}_{\mathbf{w}}]_1^T)$ 
          • forward pass for backward LSTM state, evaluating  $f_{\theta}([\mathbf{v}_{\mathbf{w}}]_1^T)$ 
      2. compute the CRF forward and backward pass
      3. compute the bidirectional LSTM-CRF backward pass for  $[\mathbf{v}_{\mathbf{w}}]_1^T$ 
          • backward pass for forward LSTM state
          • backward pass for backward LSTM state
      4. update  $[A]_{i,j}, \theta$ 
    return classifier model
  end function

```

4. Sensing Stage

4.1. Data Gathering

As presented in [64], it is challenging to retrieve all tweets during and after an event and choose them on the basis of their inherent subjectivity or authenticity of the publishing entity. As concluded in [42], there are two types of queries intended to reduce non-relevant data: (1) keyword-based queries, which search for terms and hashtags determined to be relevant; (2) geographical geo-queries, which search within a bounding box of places of interest. Our proposed sensor monitors hashtags that specifically describe a topic $\in T$. We use several geo-queries bounded to the geographical region of interest, e.g., Mexico City. Such geo-queries are aimed to retrieve tweets that contain at least one keyword-based hashtags related to the event of interest, for example for the earthquake that occurred on 19 September 2017 in Mexico City, these keyword-based hashtags are: #sismo, #sismoCDMX, #AyudaCDMX, #FuerzaMexico, #AquiNecesitamos, #derrumbe, #19s, #Voluntarios,

#ayudasismoCDMX. For this particular natural disaster, the querying terms are also complemented by well-defined urban spaces from a city [65], e.g., #derrumbe,avenida (which translates into #collapse,avenue), to guarantee that there is, as a minimum, a named place and a particular topic. Twitter characteristics, such as retweets and mentions, contribute to the widespread dissemination of a tweet reporting a location, so these features can be used as a source of temporal information [66]. To exemplify this, a query q , related to the tool developed in [20], is shown in Equation (9):

$$q = [\text{\#sismo, ayuda, avenida tlalpan}], \quad (9)$$

where q contains the following words in English: #earthquake,help,tlalpan avenue.

4.2. Spatial and Temporal Information

To be able to sense spatial information, a dataset X_s comprising tweets scraped in real time is built for the event of interest. As mentioned before, to classify every tweet into named entities, each tweet must be transformed into its word embedding representation. When several tweets are collected, X_s is fed into the classification model to obtain a series of predicted entities \hat{y} . Prior to the conversion of predictions into useful *toponyms*, words classified with 0 and *PER* tags are discarded (their presence is required in the training stage to capture entity tag transitions at the CRF output layer, but they are not needed for *toponym* identification). Furthermore, those classified as *LOC* and *ORG* are identified and joined to form a sentence, consequently creating a *toponym*, which is used to request a Google API location. Responses from Google are geocoded in JSON format to form an address with geographic coordinates. To reduce processing times, *toponyms* are appended in a set denoted by \hat{Y} and transformed into a One-Hot-Encoding vector to cluster them using the cosine similarity metric [67] (given some threshold $\alpha \in [0, 1]$). Therefore, if a requested *toponym* is similar to one that is already geocoded, it is assigned the same address and spatial information. Figure 5 depicts the proposed method for geocoding.

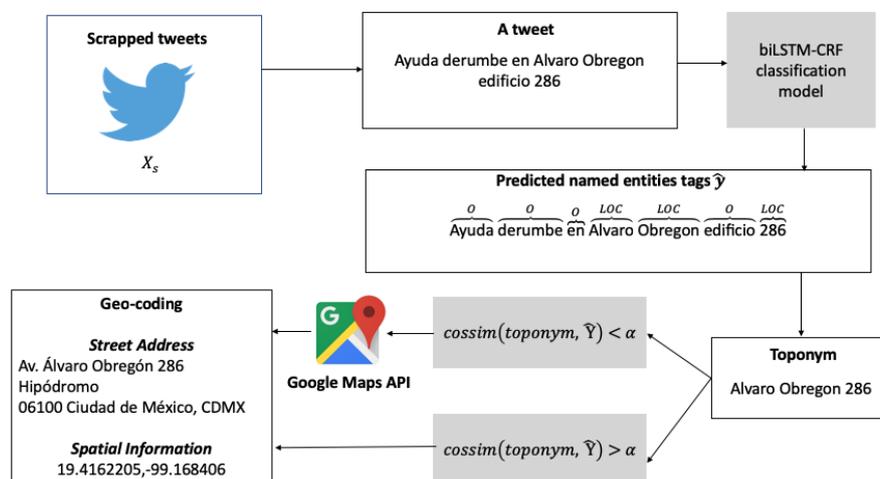


Figure 5. Toponym geocoding.

To extract temporal information, time windows are employed. In this way, the sensor can grab timestamps, ts , corresponding to the date that a tweet was created. Given the information about when a tweet was initially scraped (the first tweet naming a *toponym*), its spatial information can be foot-printed. Subsequent retweets (child nodes) originating from an initial tweet are assigned a timestamp equal to the difference between the date of creation of the parent and their current timestamp, i.e., $(ts_{parent} - ts_{children})$, $\{\forall ts_{parent} \in \hat{y}_{parent}, \forall ts_{parent} \in \hat{y}_{children}\}$. If a *toponym* is identical to others according to the threshold α , the date of creation is then calculated on the basis of those clustered by similarity. Tweets are then sorted by date of creation, from the oldest to the most recent,

i.e., $sort(\hat{y}_1 \mapsto ts_1, \dots, \hat{y}_n \mapsto ts_n)$. For practicality, a 3-day observation window with 7765 unique tweets and 14,155 retweets is applied in our case study.

4.3. Kernel Density Estimation

KDE [68] is a statistical method broadly used to graphically visualize hotspots from spatial points distributed on a two-dimensional probability density function [69–72]. KDE is used on the geocoded *toponyms* to appropriately estimate the distribution of geographic locations within the time windows previously presented in Section 4.2. By plotting with Matplotlib’s Basemap [73], it is possible to visualize geographic areas by topics $\in T$, which may include areas likely to be dangerous, plot the zones with the highest rates of missing individuals, and locate aid services via shelters. To quantify the incoming geocoded *toponyms* at a spatial point g , Equation (10) is used [72]:

$$f(g) = \gamma(g, h) = \frac{1}{Ph} \sum_{\forall \omega_i} K \left(\frac{\|g - g_{\omega_i}\|_{\ell_2}}{h} \right), \quad (10)$$

where h is the bandwidth; P is the total number of pieces of geocoded information of a topic $T \in \{\text{disaster areas, missing individuals, shelters}\}$ within the time window ω , i indexes a single geocoded *toponym* within a time window ω , K is the density function, and ℓ_2 is the vector norm.

5. Sensing Information: A Case Study of the 2017 Mexico City Earthquake

On 19 September 2017 at 1:14 p.m. CST, an earthquake with a 7.1 magnitude on the Richter scale with an epicenter in Axochiapan, Morelos, a state adjacent to Mexico City, impacted the urban infrastructure of the city and surrounding areas. Although the alarm system is efficient when epicenters occur on the Pacific Ocean coast, in the particular case of this natural disaster, the evacuations took place 11 s after the earthquake started because of the lack of sensors near the metropolitan area. It was not to be expected that Twitter users would report information related to the disaster zones. In addition to army and navy personnel, a large number of individuals took to the streets to offer humanitarian aid to people in major risk areas. Days later, a number of official and collaborative shelters were set up in churches, parks, schools, and other places to offer help to the victims. Figure 6 shows a sample of tweets sent over a 3-day observation window.

To compare the proposed sensor with the recent state of the art, a survey was taken of recent works that aimed to address natural disaster monitoring using OSN data with open and available datasets. Table 6 summarizes the works selected to be compared.

Table 6. Recent works used to compare the proposed sensor.

Titles	Natural Disaster	Dataset	Algorithms Employed	Algorithm with Overall Best Performance Metric Reported	Year
Twitter as a Lifeline: Human-annotated Twitter Corpora for NLP of Crisis-related [74]	Napa California Earthquake, USA	Publicly available on <i>Resources for Research on Crisis Informatics</i> [75]	NB, SVM, and RF with Word Embeddings	NB with 82% accuracy	2016
A linguistically-driven approach to cross-event damage assessment of natural disasters from social media messages [43] A Big Data Crisis Mapping System Based on Damage Detection and Geoparsing [76]	L’Aquila and Emilia earthquakes from 2009 to 2014, Italy	Publicly Available on <i>Project SoS</i> [77]	SVM + Word Embeddings, SVM and NLP + POS tags	SVM + Word Embeddings with 88% F1-score	2018

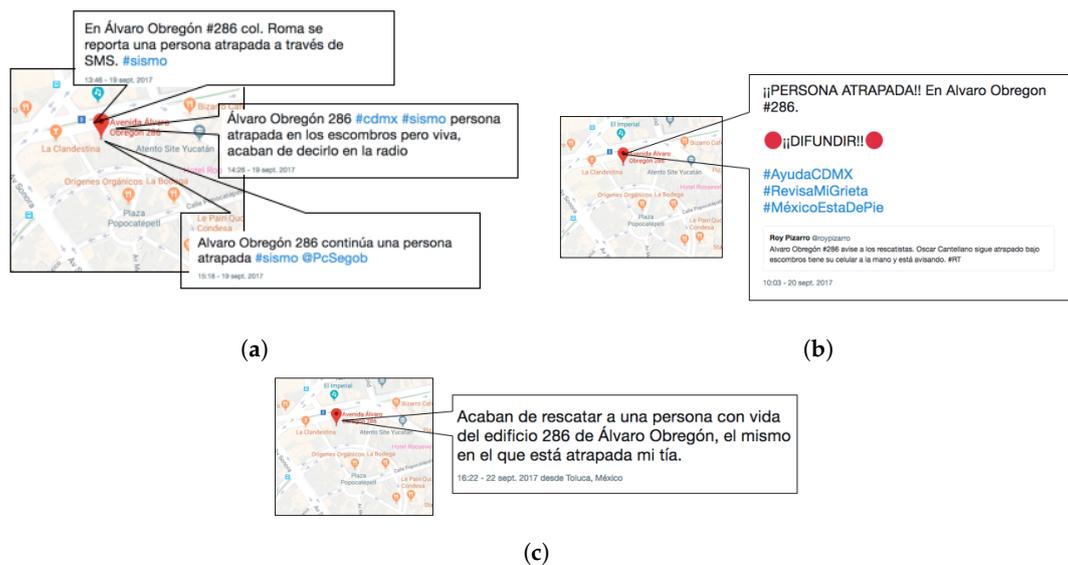


Figure 6. The first report occurs at 1:46 p.m., almost half an hour after the earthquake. The localized entity corresponds to the street *Av. Álvaro Obregón*, number 286, with geographic coordinates 19.4162205, -99.1705947 . The other classified entities are similar and ordered temporally until the last report at 4:22 p.m. on the third observation day. (a) Users first report that a person is trapped in a collapsed building; (b) a day later, users continue reporting that a person is in the rubble, and information is already disseminated in a retweet; (c) on the third day, the victim is reported as rescued.

In [74], the authors assess the impact of a natural hazard and evaluate different topics: *Caution and advice*, *Displaced people and evacuations*, *Donation needs or offers*, *Infrastructure and utilities damage*, *Injured or dead people*, *Missing, trapped or found people*, *Sympathy emotional support*, *Other useful information*, and *Not related or irrelevant*. Then, for each topic, they process tweets by removing noisy patterns, followed by tagging out-of-vocabulary words and normalizing them. Further, they weigh terms using Word2vec and use them for training three classifiers: NB, SVM, and RF. In [76], the authors employ Word Embeddings of a fixed sized and a *simple linear* kernel SVM to classify tweets into one of three topics: *Damage*, *No damage*, and *Not relevant*. To compare these two works with our methodology, datasets provided by [43,74,76] were annotated with the entity classes described in Section 3.1 using Polyglot [78], an NER tagger for multi-lingual purposes, along with other handcrafted rules. Thereafter, to evaluate classification performance, the tagged datasets and X , the corpus of tweets used in this work (Mexico City Earthquake), were trained with the pool of algorithms used in [74] and [43,76], as well as with the algorithm (biLSTM-CRF) used in our methodology. For each algorithm, it was assumed that words were preprocessed and transformed into word embedding representations. For biLSTM-CRF, only tags describing a *toponym* (LOC and ORG) are considered. The results are listed in Table 7 in terms of the *precision*, *recall*, and *F-1 score*.

As observed in Table 7, the biLSTM-CRF classifier used to build the proposed sensor performs better on average compared with the the RF, SVM, and NB algorithms. The biLSTM-CRF classifier achieves, on average, a precision = 0.85, a recall = 0.82, and an F1-score = 0.84. Even though word embeddings are used in all approaches, only the biLSTM-CRF classifier can capture the maximum contextual information in both directions of a word embedding and its transitions between NER tags at the state level (sentence), thus improving performance results.

Table 7. Results comparison

Dataset	Classifier	Named Entity Tag	Precision	Recall	F-1 Score
19 September 2017 Mexico Earthquake	biLSTM-CRF	LOC	0.83	0.76	0.80
19 September 2017 Mexico Earthquake	biLSTM-CRF	ORG	0.83	0.86	0.85
2009–2014 L’Aquila and Emilia earthquakes, Italy	biLSTM-CRF	LOC	0.84	0.84	0.84
2009–2014 L’Aquila and Emilia earthquakes, Italy	biLSTM-CRF	ORG	0.79	0.69	0.74
2014 Napa California Earthquake, USA	biLSTM-CRF	LOC	0.93	0.90	0.92
2014 Napa California Earthquake, USA	biLSTM-CRF	ORG	0.88	0.87	0.87
		Average	0.85	0.82	0.84
19 September 2017 Mexico Earthquake	RF	LOC	0.89	0.19	0.31
19 September 2017 Mexico Earthquake	RF	ORG	0.89	0.18	0.30
2009–2014 L’Aquila and Emilia earthquakes, Italy	RF	LOC	0.74	0.60	0.66
2009–2014 L’Aquila and Emilia earthquakes, Italy	RF	ORG	0.76	0.29	0.42
2014 Napa California Earthquake, USA	RF	LOC	0.60	0.25	0.35
2014 Napa California Earthquake, USA	RF	ORG	0.75	0.26	0.39
		Average	0.77	0.30	0.40
19 September 2017 Mexico Earthquake	SVM	LOC	0.76	0.48	0.59
19 September 2017 Mexico Earthquake	SVM	ORG	0.73	0.78	0.64
2009–2014 L’Aquila and Emilia earthquakes, Italy	SVM	LOC	0.75	0.57	0.65
2009–2014 L’Aquila and Emilia earthquakes, Italy	SVM	ORG	0.82	0.25	0.38
2014 Napa California Earthquake, USA	SVM	LOC	0.63	0.44	0.52
2014 Napa California Earthquake, USA	SVM	ORG	0.82	0.24	0.37
		Average	0.75	0.45	0.53
19 September 2017 Mexico Earthquake	NB	LOC	0.88	0.19	0.31
19 September 2017 Mexico Earthquake	NB	ORG	0.86	0.18	0.30
2009–2014 L’Aquila and Emilia earthquakes, Italy	NB	LOC	0.79	0.46	0.58
2009–2014 L’Aquila and Emilia earthquakes, Italy	NB	ORG	0.84	0.24	0.37
2014 Napa California Earthquake, USA	NB	LOC	0.51	0.57	0.54
2014 Napa California Earthquake, USA	NB	ORG	0.78	0.26	0.39
		Average	0.70	0.47	0.42

Visualizing the Social Dynamics via KDE

Figure 7a–c depict the hotspots obtained by KDE estimations from the geocoded *toponyms* over a span of three days. These hotspots allow visualizing areas with the highest concentration of tweets reporting a specific topic and naming a *toponym*; i.e., $T \in \{\text{disaster areas, missing individuals, shelters}\}$. To validate these results, these hotmaps are compared with two collaborative maps populated with official data verified by Google and the Mexican government (publicly available as *Mapeo Verificado19s* [79]). The information contained in *Mapeo Verificado19s* maps is divided into the following categories:

- **Official Damages:** includes collapsed buildings, major and minor risks, and wall collapses.
- **Official Shelters:** official government assistance and aid.

- **Collaborative Shelters:** non-official collaborative assistance and aid.

In addition, sources of information that contributed in a collaborative way to the population of maps during and after the earthquake are listed below:

- **Mexico City's Monitor System:** includes major risks, collapsed buildings, and gas hazards.
- **Harvard-Massachusetts Institute of Technology (MIT):** collaborative data gathered from social media sources.

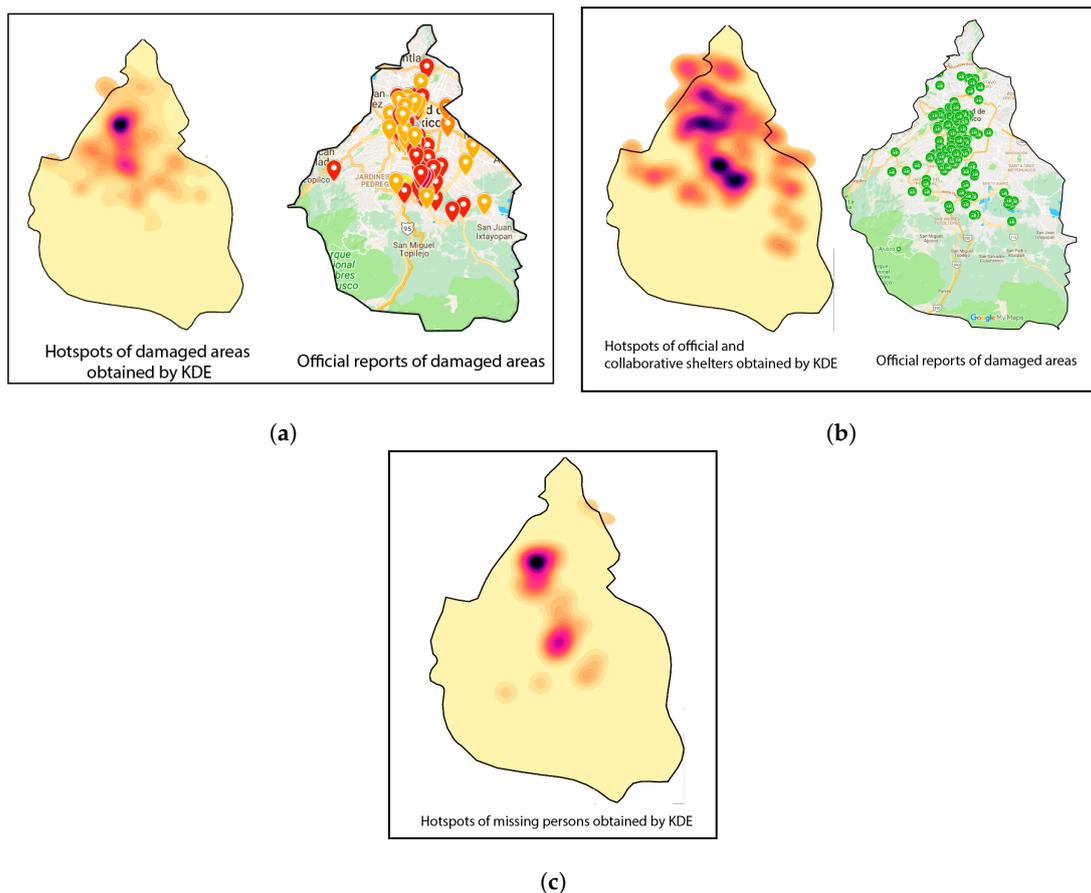


Figure 7. Hotspots maps obtained by applying KDE to the spatial information extracted from data collected over a 3-day window. (a) The hotspot map of the estimated spatial locations related to damages and collapses and official reports. (b) The hotspot map of estimated spatial locations related to official and collaborative shelters and official reports. (c) The hotspot map of estimated spatial locations related to missing persons (there are no official reports of missing persons).

It is important to emphasize that official and collaborative maps, e.g., *Mapeo Verificado19s*, neither allow for determining the spatial density of the topic of interest nor account for missing persons. This can be a crucial disadvantage in cases where it is necessary to examine the dynamics and evolution of an event of interest on the basis of incoming reports. The authenticity of *toponyms* is tested by searching official addresses published by Mexico's federal government [80]. This information can be collected only after civil protection units verify the geographical areas of the disaster and issue an official statement. Unfortunately, there were no official data for this natural disaster related to aid, shelter, and missing persons. The proposed sensor has then the potential to assist in estimating in real-time the geographical regions with the largest density of tweets associated with a specific topic of interest, enabling information to be disseminated without subjecting responders to the risks associated with on-site verification. Table 8 lists the most common geocoded *toponyms* transformed into Google

API addresses found by our sensor. These locations have also been officially declared as disaster areas by Mexico's federal government.

Table 8. Geocoded addresses and coordinates found by the sensor and officially declared as disaster areas.

	Geocoded Address	Geocoded Coordinates	Tweets	Retweets
1	Rancho Tamboreo & Calz de las Brujas, Nueva Oriental Coapa, 14300 Ciudad de México, CDMX	19.2965695, −99.1328497	135	368
2	Calz. de Tlalpan 20, Conjunto Urbano Tlalpan, 04400 Ciudad de México, CDMX	19.3385929, −99.1446581	126	331
3	Av. Álvaro Obregón 286 Hipódromo 06100 Ciudad de México, CDMX	19.4162255, −99.170594	112	250
4	Amsterdam 25, Hipódromo, 06100 Ciudad de México, CDMX	19.4158929, −99.1701461	109	204
5	Calle Torreón & Viad. Miguel Alemán, Piedad Narvarte, 06760 Ciudad de México, CDMX	19.4025116, −99.1634792	104	237
6	Edimburgo & Escocia, Col del Valle Centro, 03100 Ciudad de México, CDMX	19.3875319, −99.1656197	103	228
7	Amsterdam & Calle Laredo, Hipódromo, 06100 Ciudad de México, CDMX	19.4129041, −99.1730674	97	143
8	Av. Álvaro Obregón 284, Hipódromo, 06100 Ciudad de México, CDMX	19.4162562, −99.1704433	96	127
9	Coahuila 286 Hipódromo, 06700 Ciudad de México, CDMX	19.410391, −99.1685889	94	164
10	Simón Bolívar 190, Obrera, 06800 Ciudad de México, CDMX	19.4221723, −99.1422295	95	131
11	Petén & Gral. Emiliano Zapata, Sta Cruz Atoyac, 03320 Ciudad de México, CDMX	19.3665055, −99.1591011	92	199
12	Puebla 282 Roma Nte. 06700 Ciudad de México, CDMX	19.4211364, −99.1714281	92	216
13	Calle Salamanca 107, Roma Nte., 06700 Ciudad de México, CDMX	19.4172303, −99.1714257	91	139
14	Balsas 18 sineo, Miravalle 03580 Ciudad de México, CDMX	19.3605422, −99.1424208	88	215
15	Escocia & Calle Gabriel Mancera, Col del Valle Centro, 03100 Ciudad de México, CDMX	19.3876749, −99.1661223	87	220
16	Calz. de Tlalpan 2050, Campestre Churubusco, 04200 Ciudad de México, CDMX	19.3429739, −99.1434801	74	155
17	Calle Querétaro & Medellín, Roma Nte. 06700 Ciudad de México, CDMX	19.413905, −99.1672667	73	211
18	Av Sonora 149, Hipódromo, 06100 Ciudad de México, CDMX	19.4145946, −99.1714381	70	237
19	Calle Concepción Beistegui & Calle Yacatas, Narvarte Poniente 03020 Ciudad de México, CDMX	19.3873507, −99.1582722	69	178
20	Galicia Niños Héroe, Ciudad de México, CDMX	19.3886011, −99.1482661	69	111
21	Calle Enrique Rebsamen & La Morena Narvarte Poniente, 03020 Ciudad de México, CDMX	19.3985479, −99.1609147	61	97
22	Rancho Vista Hermosa & Rancho de Los Arcos, Parque Alameda del Sur 04929 Ciudad de México, CDMX	19.3069132, −99.124864	54	128
23	Bretaña & Irolo, Zacahuitzco, 03550 Ciudad de México, CDMX	19.3731238, −99.1398383	50	133
24	Gral. Emiliano Zapata 51 Portales Nte, Ciudad de México, CDMX	19.3642598, −99.1446719	47	131

Table 8. Cont.

	Geocoded Address	Geocoded Coordinates	Tweets	Retweets
25	Saratoga 714, Portales Sur, 03303 Ciudad de México, CDMX	19.3649279, −99.1540524	43	94
26	Sierravista & Calle Riobamba Lindavista Nte. 07300 Ciudad de México, CDMX	19.4940873, −99.1265294	41	117
27	Calle Salvador Díaz Mironn Sta María la Ribera Ciudad de México, CDMX	19.4492376, −99.1620973	40	85
28	Av. las Trancas 40 Narciso Mendoza 14390 U. Hab. Narciso Mendoza Super 6 Coapa, CDMX	19.292755, −99.125329	38	72
29	Calz. de la Viga 1756, Héroes de Churubusco, 09090 Ciudad de México, CDMX	19.3612758, −99.1240497	37	101
30	Avenida Santa Ana 300, Ex-Ejido de San Francisco Culhuacan, 04470 Ciudad de México, CDMX	19.3296075, −99.1272789	37	99
31	Coquimbo 07300 Ciudad de México, CDMX	19.4899307, −99.1281605	31	91
32	Calle Puente 222, San Bartolo el Chico, 14380 Ciudad de México, CDMX	19.2833487, −99.1373406	24	79
33	Paseo Galias 47 Lomas Estrella 2da Secc, 09890, Ciudad de México, CDMX	19.3205935, −99.0995659	23	56
34	Vicente Guerrero 40, San Gregorio Atlapulco, 16600, Ciudad de México, CDMX	19.2522187, −99.0614642	16	67
35	Av. México, San Gregorio Atlapulco, 16600 Ciudad de México, CDMX	19.2531664, −99.0513852,	15	74
36	Xochimilco-tulyehualco 191, Xochimilco, 16500, Ciudad de México, CDMX	19.2468579, −99.0835714	13	46

6. Conclusions

In this work, a methodology that uses Twitter as a social sensor is proposed. This is accomplished by employing an information sequential extraction procedure known as Named Entity Recognition (NER), which aims to describe mentioned entities, such as places, persons, and organizations. The methodology considers the semantic, morphological, and contextual information about each word composing a tweet and its surrounding context, thus allowing to properly identify a named place (*toponym*). To achieve this, words are tokenized and transformed into word embeddings to represent them as vectors with rich syntactic and semantic relationships that are established by neighboring words. To ensure that a high classification accuracy of the sequential data is achieved with out heavily relying on handcrafted feature extraction techniques, a Recurrent Neural Network variant, i.e., a Bidirectional Long Short-Term Memory (biLSTM) network, is used. Specifically, the biLSTM network deals with long-distance dependencies, which feed-forward algorithms, such as NB, SVM, and RF, cannot handle. This is achieved by considering contextual information in both directions of a word in a tweet. By using a CRF output layer with the biLSTM network, NER tag transitions over the word embeddings are accounted for.

In the presented case study, geo-queries related to the earthquake of 19 September 2017 in Mexico City were used to retrieve tweets with specific keyword-based hashtags. After classifying Tweets with NER tags and joining them to form useful *toponyms*, these *toponyms* were geocoded in terms of addresses and latitude and longitude coordinates by means of Google's API. Finally, a KDE algorithm was computed to visualize the spatial density of geocoded *toponyms* from topics related to disaster areas, missing individuals, and shelters. Our results show that addresses and coordinates obtained by our methodology coincide with the ones reported by civil protection units and with official data from Mexico's federal government. Collaborating with the government and civil organizations to improve

the timely detection of disaster areas, finding missing individuals, and locating shelters in real-time by using our proposed methodology is part of our future work.

Author Contributions: All authors contributed equally to this work.

Funding: This project has received funding from the European Union’s Horizon 2020 research and innovation programme under Grant agreement No. 700326. Website: <http://ramses2020.eu>. 

Acknowledgments: The authors thank the National Science and Technology Council of Mexico (CONACyT), and the Instituto Politécnico Nacional for the financial support for this research.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Ekström, G.; Dziewonski, A.M.; Steim, J.M. Single station CMT; Application to the Michoacan, Mexico, earthquake of September 19, 1985. *Geophys. Res. Lett.* **1986**, *13*, 173–176. [[CrossRef](#)]
- Gao, H.; Barbier, G.; Goolsby, R.; Zeng, D. *Harnessing the Crowdsourcing Power of Social Media for Disaster Relief*; Arizona State Univ Tempe: Tempe, AZ, USA, 2011.
- Teets, J.C. Post-earthquake relief and reconstruction efforts: The emergence of civil society in China? *China Q.* **2009**, *198*, 330–347. [[CrossRef](#)]
- Smith, P.C.; Simpson, D.M. Technology and communications in an urban crisis: The role of mobile communications systems in disasters. *J. Urban Technol.* **2009**, *16*, 133–149. [[CrossRef](#)]
- Heinzelman, J.; Waters, C. *Crowdsourcing Crisis Information in Disaster-Affected Haiti*; US Institute of Peace: Washington, DC, USA, 2010.
- Historias de WhatsApp que Salvaron Vidas Tras el Sismo en México. Available online: <http://www.eluniv.ersal.com.mx/techbit/historias-de-whatsapp-que-salvaron-vidas-tras-el-sismo-en-mexico> (accessed on 1 July 2018).
- Dhillon, H.S.; Huang, H.; Viswanathan, H. Wide-area wireless communication challenges for the Internet of Things. *IEEE Commun. Mag.* **2017**, *55*, 168–174. [[CrossRef](#)]
- Hayashi, N.J.; Ott, E.S., IV; Tsang, A.Y.; Fukuda, M.; Wascovich, D.; Quoc, M. Multimedia Sharing in Social Networks for Mobile Devices. U.S. Patent No. 8,046,411, 25 October 2011.
- Kaplan, A.M.; Haenlein, M. Users of the world, unite! The challenges and opportunities of Social Media. *Bus. Horiz.* **2010**, *53*, 59–68. [[CrossRef](#)]
- García-Palomares, J.C.; Salas-Olmedo, M.H.; Moya-Gómez, B.; Condeco-Melhorado, A.; Gutierrez, J. City dynamics through Twitter: Relationships between land use and spatiotemporal demographics. *Cities* **2018**, *72*, 310–319. [[CrossRef](#)]
- Sagl, G.; Resch, B.; Hawelka, B.; Beinat, E. From social sensor data to collective human behaviour patterns: Analysing and visualising spatio-temporal dynamics in urban environments. In Proceedings of the GI_Forum 2012: Geovisualization, Society and LearningGI-Forum, Salzburg, Austria, 3–6 July 2012; Herbert Wichmann Verlag: Berlin, Germany, 2012; pp. 54–63.
- Aggarwal, C.C.; Abdelzaher, T. Social sensing. In *Managing and Mining Sensor Data*; Springer: Boston, MA, USA, 2013; pp. 237–297.
- Aggarwal, C.C. (Ed.) *Managing and Mining Sensor Data*; Springer Science & Business Media: Berlin, Germany, 2013.
- Abdelzaher, T.; Anokwa, Y.; Boda, P.; Burke, J.; Estrin, D.; Guibas, L.; Kansal, A.; Madden, S.; Reich, J. Mobiscopes for human spaces. *IEEE Pervasive Comput.* **2007**, *6*, 20–29. [[CrossRef](#)]
- Xu, Z.; Mei, L.; Choo, K.K.R.; Lv, Z.; Hu, C.; Luo, X.; Liu, Y. Mobile crowd sensing of human-like intelligence using social sensors: A survey. *Neurocomputing* **2018**, *279*, 3–10. [[CrossRef](#)]
- Wang, R.Q.; Mao, H.; Wang, Y.; Rae, C.; Shaw, W. Hyper-resolution monitoring of urban flooding with social media and crowdsourcing data. *Comput. Geosci.* **2018**, *111*, 139–147. [[CrossRef](#)]
- Kursuncu, U.; Gaur, M.; Lokala, U.; Thirunarayan, K.; Sheth, A.; Arpinar, I.B. Predictive Analysis on Twitter: Techniques and Applications. In *Emerging Research Challenges and Opportunities in Computational Social Network Analysis and Mining*; Springer: Cham, Switzerland, 2019; pp. 67–104.
- Gaber, I. Twitter: A useful tool for studying elections? *Convergence* **2017**, *23*, 603–626. [[CrossRef](#)]

19. Pond, P.; Lewis, J. Riots and Twitter: Connective politics, social media and framing discourses in the digital public sphere. *Inf. Commun. Soc.* **2019**, *22*, 213–231. [[CrossRef](#)]
20. Hernandez-Suarez, A.; Sanchez-Perez, G.; Toscano-Medina, K.; Martinez-Hernandez, V.; Perez-Meana, H.; Olivares-Mercado, J.; Sanchez, V. Social Sentiment Sensor in Twitter for Predicting Cyber-Attacks Using ℓ_1 Regularization. *Sensors* **2018**, *18*, 1380. [[CrossRef](#)]
21. Hart, A.G.; Carpenter, W.S.; Hlustik-Smith, E.; Reed, M.; Goodenough, A.E. Testing the potential of Twitter mining methods for data acquisition: Evaluating novel opportunities for ecological research in multiple taxa. *Methods Ecol. Evol.* **2018**, *9*, 2194–2205. [[CrossRef](#)]
22. Lee, K.; Ganti, R.; Srivatsa, M.; Mohapatra, P. Spatio-temporal provenance: Identifying location information from unstructured text. In Proceedings of the 2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops), San Diego, CA, USA, 18–22 March 2013; pp. 499–504.
23. Li, R.; Lei, K.H.; Khadiwala, R.; Chang, K.C.C. Tedas: A twitter-based event detection and analysis system. In Proceedings of the 2012 IEEE 28th International Conference on Data Engineering (ICDE), Washington, DC, USA, 1–5 April 2012; pp. 1273–1276.
24. Feng, W.; Zhang, C.; Zhang, W.; Han, J.; Wang, J.; Aggarwal, C.; Huang, J. STREAMCUBE: Hierarchical spatio-temporal hashtag clustering for event exploration over the twitter stream. In Proceedings of the 2015 IEEE 31st International Conference on Data Engineering (ICDE), Seoul, Korea, 13–17 April 2015; pp. 1561–1572.
25. Sisco, M.R.; Bosetti, V.; Weber, E.U. When do extreme weather events generate attention to climate change? *Clim. Chang.* **2017**, *143*, 227–241. [[CrossRef](#)]
26. Nadeau, D.; Sekine, S. A survey of named entity recognition and classification. *Lingvist. Investig.* **2007**, *30*, 3–26.
27. Bontcheva, K.; Derczynski, L.; Roberts, I. Crowdsourcing named entity recognition and entity linking corpora. In *Handbook of Linguistic Annotation*; Springer: Dordrecht, The Netherlands, 2017; pp. 875–892.
28. Jeon, Y.; Cho, C.; Seo, J.; Kwon, K.; Park, H.; Chung, I.J. Rule-Based Topic Trend Analysis by Using Data Mining Techniques. In *Advanced Multimedia and Ubiquitous Engineering*; Springer: Singapore, 2017; pp. 466–473.
29. Mikolov, T.; Chen, K.; Corrado, G.; Dean, J. Efficient estimation of word representations in vector space. *arXiv* **2013**, arXiv:1301.3781.
30. Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G.S.; Dean, J. Distributed representations of words and phrases and their compositionality. In Proceedings of the Advances in Neural Information Processing Systems 26 (NIPS 2013), Lake Tahoe, NV, USA, 5–10 December 2013; pp. 3111–3119.
31. Hochreiter, S.; Schmidhuber, J. Long short-term memory. *Neural Comput.* **1997**, *9*, 1735–1780. [[CrossRef](#)] [[PubMed](#)]
32. Nio, L.; Murakami, K. Japanese Sentiment Classification Using Bidirectional Long Short-Term Memory Recurrent Neural Network. In Proceedings of the 24th Annual Meeting Association for Natural Language Processing, Okayama, Japan, 12–16 March 2018.
33. Lafferty, J.; McCallum, A.; Pereira, F.C. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In Proceedings of the 18th International Conference on Machine Learning 2001 (ICML 2001), Williams College, Williamstown, MA, USA, 28 June–1 July 2001; pp. 282–289.
34. Silverman, B.W. *Density Estimation for Statistics and Data Analysis*; Routledge: Abingdon, UK, 2018.
35. Kongthon, A.; Haruechaiyasak, C.; Pailai, J.; Kongyoung, S. The role of Twitter during a natural disaster: Case study of 2011 Thai Flood. In Proceedings of the PICMET'12 Technology Management for Emerging Technologies (PICMET), Vancouver, BC, Canada, 29 July–2 August 2012; pp. 2227–2232.
36. Sachdeva, S.; McCaffrey, S. Using Social Media to Predict Air Pollution during California Wildfires. In Proceedings of the ACM 9th International Conference on Social Media and Society, Copenhagen, Denmark, 18–20 July 2018; pp. 365–369.
37. Hughes, A.L.; Palen, L. Twitter adoption and use in mass convergence and emergency events. *Int. J. Emerg. Manag.* **2009**, *6*, 248–260. [[CrossRef](#)]
38. Earle, P.S.; Bowden, D.C.; Guy, M. Twitter earthquake detection: Earthquake monitoring in a social world. *Ann. Geophys.* **2012**, *54*. [[CrossRef](#)]

39. Sakaki, T.; Okazaki, M.; Matsuo, Y. Earthquake shakes Twitter users: Real-time event detection by social sensors. In Proceedings of the ACM 19th International Conference on World Wide Web, Raleigh, NC, USA, 26–30 April 2010; pp. 851–860.
40. Finch, K.C.; Snook, K.R.; Duke, C.H.; Fu, K.W.; Tse, Z.T.H.; Adhikari, A.; Fung, I.C.H. Public health implications of social media use during natural disasters, environmental disasters, and other environmental concerns. *Nat. Hazards* **2016**, *83*, 729–760. [[CrossRef](#)]
41. Middleton, S.E.; Middleton, L.; Modafferi, S. Real-time crisis mapping of natural disasters using social media. *IEEE Intell. Syst.* **2014**, *29*, 9–17. [[CrossRef](#)]
42. Ashktorab, Z.; Brown, C.; Nandi, M.; Culotta, A. Tweedr: Mining twitter to inform disaster response. In Proceedings of the 11th International Conference on Information Systems for Crisis Response and Management, University Park, PA, USA, 18–21 May 2014.
43. Cresci, S.; Tesconi, M.; Cimino, A.; Dell’Orletta, F. A linguistically-driven approach to cross-event damage assessment of natural disasters from social media messages. In Proceedings of the ACM 24th International Conference on World Wide Web, Florence, Italy, 18–22 May 2015; pp. 1195–1200.
44. Resch, B.; Usländer, F.; Havas, C. Combining machine-learning topic models and spatio-temporal analysis of social media data for disaster footprint and damage assessment. *Cartogr. Geogr. Inf. Sci.* **2018**, *45*, 362–376.
45. Matheson, D. The performance of publicness in social media: Tracing patterns in tweets after a disaster. *Media Cult. Soc.* **2018**, *40*, 584–599. [[CrossRef](#)]
46. Gruebner, O.; Lowe, S.; Sykora, M.; Shankardass, K.; Subramanian, S.V.; Galea, S. Spatio-temporal distribution of negative emotions in New York City after a natural disaster as seen in social media. *Int. J. Environ. Res. Public Health* **2018**, *15*, 2275. [[CrossRef](#)]
47. Zhang, W.; Yoshida, T.; Tang, X. A comparative study of TF* IDF, LSI and multi-words for text classification. *Expert Syst. Appl.* **2011**, *38*, 2758–2765. [[CrossRef](#)]
48. Schmolz, H. *Anaphora Resolution and Text Retrieval: A Linguistic Analysis of Hypertexts*; Walter de Gruyter GmbH & Co KG: Berlin, Germany, 2015.
49. Sravani, L.; Reddy, A.S.; Thara, S. A Comparison Study of Word Embedding for Detecting Named Entities of Code-Mixed Data in Indian Language. In Proceedings of the 2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI), Bangalore, India, 19–22 September 2018; pp. 2375–2381.
50. Young, T.; Hazarika, D.; Poria, S.; Cambria, E. Recent trends in deep learning based natural language processing. *IEEE Comput. Intell. Mag.* **2018**, *13*, 55–75. [[CrossRef](#)]
51. Hernandez-Suarez, A.; Sanchez-Perez, G.; Martinez-Hernandez, V.; Olivares Mercado, J. Can Twitter API Be Bypassed? A New Methodology for Collecting Chronological Information Without Restrictions. In Proceedings of the 17th in New Trends in Intelligent Software Methodologies, Tools and Techniques International Conference, Granada, Spain, 26–28 September 2018.
52. Erol, M.H.; Bulut, F. Real-time application of travelling salesman problem using Google Maps API. In Proceedings of the IEEE Electric Electronics, Computer Science, Biomedical Engineerings’ Meeting (EBBT), Istanbul, Turkey, 20–21 April 2017; pp. 1–5.
53. Ratnov, L.; Roth, D. Design challenges and misconceptions in named entity recognition. In Proceedings of the Thirteenth Conference on Computational Natural Language Learning, Shared Task Boulder, CO, USA, 4–5 June 2009; pp. 147–155.
54. Chieu, H.L.; Ng, H.T. Named entity recognition: A maximum entropy approach using global information. In Proceedings of the 19th International Conference on Computational Linguistics, Taipei, Taiwan, 24 August–1 September 2002; Volume 1; p. 786.
55. Tjong Kim Sang, E.F.; De Meulder, F. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL, Edmonton, AB, Canada, 31 May–1 June 2003; pp. 142–147.
56. Turian, J.; Ratnov, L.; Bengio, Y. Word representations: A simple and general method for semi-supervised learning. In Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, Uppsala, Sweden, 11–16 July 2010; pp. 384–394.
57. Liu, X.; Zhang, S.; Wei, F.; Zhou, M. Recognizing named entities in tweets. In Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Portland, OR, USA, 19–20 June 2011; pp. 359–367.
58. Pang, B.; Lee, L. Opinion mining and sentiment analysis. *Found. Trends Inf. Retr.* **2008**, *2*, 11–35. [[CrossRef](#)]

59. Goldberg, Y.; Levy, O. word2vec Explained: Deriving Mikolov et al.'s negative-sampling word-embedding method. *arXiv* **2014**, arXiv:1402.3722.
60. Al-Smadi, M.; Talafha, B.; Al-Ayyoub, M.; Jararweh, Y. Using long short-term memory deep neural networks for aspect-based sentiment analysis of Arabic reviews. *Int. J. Mach. Learn. Cybern.* **2018**, *2*, 1–13. [CrossRef]
61. Greenberg, N.; Bansal, T.; Verga, P.; McCallum, A. Marginal Likelihood Training of BiLSTM-CRF for Biomedical Named Entity Recognition from Disjoint Label Sets. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, 31 October–4 November 2018; pp. 2824–2829.
62. Do, H.; Than, K.; Larmande, P. Evaluating Named-Entity Recognition approaches in plant molecular biology. In Proceedings of the International Conference on Multi-Disciplinary Trends in Artificial Intelligence, Hanoi, Vietnam, 18–20 November 2018.
63. Xu, J.; He, H.; Sun, X.; Ren, X.; Li, S. Cross-Domain and Semisupervised Named Entity Recognition in Chinese Social Media: A Unified Model. *IEEE/ACM Trans. Audio Speech Lang. Process.* **2018**, *26*, 2142–2152. [CrossRef]
64. Bruns, A.; Liang, Y.E. Tools and methods for capturing Twitter data during natural disasters. *First Monday* **2012**, *17*, 4. [CrossRef]
65. Krier, R.; Rowe, C. *Urban Space*; Academy Editions: London, UK, 1979.
66. Spiro, E.; Irvine, C.; DuBois, C.; Butts, C. Waiting for a retweet: Modeling waiting times in information propagation. In Proceedings of the 2012 NIPS Workshop of Social Networks and Social Media Conference, Evanston, IL, USA, 7–8 December 2012; Volume 12. Available online: <http://snap.stanford.edu/social2012/papers/spiro-dubois-butts.pdf> (accessed on 18 June 2018).
67. Steinbach, M.; Karypis, G.; Kumar, V. A comparison of document clustering techniques. In Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Boston, MA, USA, 20–23 August 2000; Volume 400, pp. 525–526.
68. Sheather, S.J.; Jones, M.C. A reliable data-based bandwidth selection method for kernel density estimation. *J. R. Stat. Soc. Ser. B Methodol.* **1991**, *53*, 683–690. [CrossRef]
69. Li, L.; Goodchild, M.F.; Xu, B. Spatial, temporal, and socioeconomic patterns in the use of Twitter and Flickr. *Cartogr. Geogr. Inf. Sci.* **2013**, *40*, 61–77. [CrossRef]
70. Sims, K.M.; Weber, E.M.; Bhaduri, B.L.; Thakur, G.S.; Resseguie, D.R. Application of social media data to high-resolution mapping of a special event population. In *Advances in Geocomputation*; Springer: Cham, Switzerland, 2017; pp. 67–74.
71. Huang, H.; Dong, Y.; Tang, J.; Yang, H.; Chawla, N.V.; Fu, X. Will Triadic Closure Strengthen Ties in Social Networks? *ACM Trans. Knowl. Discov. Data* **2018**, *12*, 30. [CrossRef]
72. Gerber, M.S. Predicting crime using Twitter and kernel density estimation. *Decis. Support Syst.* **2014**, *61*, 115–125. [CrossRef]
73. Hunter, J.D. Matplotlib: A 2D graphics environment. *Comput. Sci. Eng.* **2007**, *9*, 90–95. [CrossRef]
74. Imran, M.; Mitra, P.; Castillo, C. Twitter as a lifeline: Human-annotated twitter corpora for NLP of crisis-related messages. *arXiv* **2016**, arXiv:1605.05894.
75. CrisisNLP. Available online: <https://crisisnlp.qcri.org/> (accessed on 1 February 2019).
76. Avvenuti, M.; Cresci, S.; Del Vigna, F.; Fagni, T.; Tesconi, M. CrisMap: A Big Data Crisis Mapping System Based on Damage Detection and Geoparsing. *Inf. Syst. Front.* **2018**. [CrossRef]
77. Project SOS. Available online: <http://socialsensing.it/en/datasets> (accessed on 1 February 2019).
78. Al-Rfou, R.; Perozzi, B.; Skiena, S. Polyglot: Distributed word representations for multilingual nlp. *arXiv* **2013**, arXiv:1307.1662.
79. Sismo Verificado 19s. Available online: <http://google.org/crisismap/a/gmail.com/v19s> (accessed on 5 July 2018).
80. Daños y Derrumbes en Edificios y Estructuras por el Sismo 19-S. Available online: <https://datos.gob.mx/busca/dataset/danos-y-derrumbes-en-edificios-y-estructuras-por-el-sismo-19-s> (accessed on 5 July 2018).

